LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-JRNL-420867

Extreme Scale Computing to Secure the Nation

D. L. Brown, J. McGraw, D. Nikkel, D. Pilkington, J. R. Johnson, D. Frincke

December 1, 2009

SciDAC Review

Disclaimer

Extreme Scale Computing to Secure the Nation

Contributors: David L. Brown, Jim McGraw, Daniel J. Nikkel, Des Pilkington, Lawrence Livermore National Laboratory; John R. Johnson and Deb Frincke, Pacific Northwest National Laboratory.

What do you do when an aging critical component in one of your few remaining nuclear stockpile weapons must be replaced, and for health and safety reasons it must be replaced by a different material? Either you must withdraw one of the few remaining nuclear systems, or you must be able to confirm to the President, that the weapon with the corrected component will perform as required and still be safe and secure. Additionally, without the benefit of a confirmatory nuclear experiment or the involvement of staff with nuclear test experience, you must be able to demonstrate confidence in your recommendation! Since the beginning of the US nuclear test moratorium in the early 1990's , this is the scenario that the Department of Energy (DOE), has been preparing to face, and must be able to address by the end of the next decade.

The DOE response to this challenge was to begin to transition the role of computer simulation from an aid for the experienced nuclear weapon designer, to a modern predictive capability that would not only provide a quantified assessment of a nuclear device safety and performance, but would be much more widely applicable to other issues of National Security. As will be demonstrated in this article, achievement of this predictive capability goal will require the power of exascale computing supported by the development and deployment of computers and the attendant programming models and software environments.

The DOE initiated the Science-Based Stockpile Stewardship program (SSP) in response to the Fiscal Year 1994 National Defense Authorization Act, which requires, "in the absence of nuclear testing, a program to:

1. Support a focused, multifaceted program to increase the understanding of the enduring stockpile;
2. Predict, detect, and evaluate potential problems of the aging of the stockpile;
3. Refurbish and re-manufacture weapons and components, as required; and
4. Maintain the science and engineering institutions needed to support the nation's nuclear deterrent, now and in the future".

A major component of this endeavor was the Accelerated Strategic Computing Initiative (ASCI), which took on the challenge of developing the three dimensional (3D) multi-physics computer programs ("codes"), and the supercomputer environments that would be needed to run these codes. It would be these codes that would be used to assess the present, and future, stockpile *without* reliance upon calibration against past, or future, nuclear test data. However it was not only the nuclear stockpile that was aging, so too were the nuclear weapon designers, who

even today, are the repository for the vast majority of the knowledge relating to the sensitivities and subtleties of the complex performance and safety issues relating to nuclear weapons. It was a race against time!

Having successfully demonstrated the feasibility of massively parallel, high resolution, 3D computing , the Initiative aspect of ASCI has ended and ASCI has been replaced by the Advanced Simulation and Computing (ASC) program within the National Nuclear Security Administration (NNSA) within the DOE. ASC now has the responsibility to deliver the Nuclear Security predictive simulation capability that will enable the certification of the nuclear stockpile into the future. (sidebar: "History of ASC Computing Capability"). **INSERT BROWNSIDEBAR1.** One arm of ASC's roadmap to meet this goal requires the development of computing capabilities at the exascale level -- capabilities that are a billion times more capable than those that were used in the last days of nuclear testing.

So what drives such an enormous increase in computing requirements? The answer is deceptively simple: significant improvements in geometric fidelity (e.g. representation of initial conditions), in physics fidelity (e.g. fundamental physics approximations & equations), and in numerical fidelity (e.g. resolution & uncertainty assessments). In other words, predictive capability via exascale computing is driven by the need to solve the right problem, using the right equations, to the right level of accuracy.

The rest of this article will elaborate some of these needs, in particular areas that are of key importance for National Security.

Exascale Needs for Stockpile Stewardship

Detailed descriptions of the simulations required to support Stockpile Stewardship cannot be made in this article. However three examples are provided below (3D, physics fidelity, and uncertainty quantification), that cover the types of computing requirements that are needed, and outline how the path to delivering a predictive simulation capability will require exascale computing.

3D. The nuclear stockpile was initially designed using one- (1D) and two-dimensional (2D) codes. These codes were strongly calibrated to the underground test (UGT) base and while sufficient to "interpolate", the test results were generally not adequate to extrapolate significantly beyond the conditions of the original tests. An NNSA aim has been to remove these UGT calibrations by improving the quality of the weapon simulations. Much progress has been made but there remain a small number of high-level calibrations ("knobs") in the weapon simulation methodology. An integrated NNSA program aimed at removing these knobs has now been established, and is called the Predictive Capability Framework (PCF).

If one considers hydrodynamic instabilities, material variations that occur due to aging, or the precise initial conditions after the engineering and manufacturing

processes have been completed, then the phenomena of interest are inherently three-dimensional (3D). The first step of ASCI was therefore to improve the geometric fidelity of nuclear weapon simulations by being able to represent the three dimensions of reality. For 3D time dependent problems, computing requirements scale approximately as (mesh size)^4, 3 powers for the 3 spatial dimensions and another for the time dimension. (sidebar: "How a computational simulation is created") . **INSERT BROWNSIDEBAR2** This implies that a need to reduce the computational grid spacing by an order of magnitude will increase computational requirements by 10,000!

In 2006, during the testing of the 100 teraflop class supercomputer, ASC Purple, some remarkable Stockpile Stewardship simulations supporting the PCF were performed. These ground breaking 2D simulations were able to use a computational grid spacing an order of magnitude smaller than had ever been used before. These simulations identified new, previously unresolved, phenomena and it is believed that these phenomena will be helpful in removing one of the remaining UGT calibrations that is still routinely used in nuclear weapon simulations. Although these simulations were intriguing, the phenomena observed were inherently 3D. Performing full 3D simulations of weapon systems at this higher resolution is not possible with ASC Purple , which can still only perform calculations at "normal" resolution in 3D . Running simulations at this new, order of magnitude finer, numeric fidelity, demands about 10,000 times the capability of Purple, i.e. exascale computing resources.

Physics Fidelity. The second example of increased computing requirements relates to the need to enhance the physics fidelity of the SSP simulations. A fundamental aspect of a nuclear weapon simulation is the material property data, used to describe the behavior of a material. Material behaves differently depending upon the conditions it is put under and a wealth of data is needed to represent all the possible conditions a nuclear weapon experiences during its lifecycle. Simple examples of material property data are nuclear cross sections, radiative opacities, hydrodynamic equations of state (EOS), material strength properties, etc. The quality and accuracy of this material data will be a critical ingredient in establishing a predictive capability. Simply put, a nuclear weapon simulation will only be as good as the material property data that it uses.

Historically, where data could not be obtained by experiments, it could quite often be inferred from UGTs, or at least the simulated integrated behavior could be calibrated against UGTs. Without UGTs, this data needs to be derived from other sources. This has led to an enhancement in the NNSA non-nuclear experimental capabilities (e.g. The Dual-Axis Radiographic Hydrodynamic Test Facility at LANL, and the National Ignition Facility at LLNL), but with the continuing improvement of simulation capability, simulations at high resolution are now another potential path to discovering material data. (sidebar: "The role of computational resolution in determining transition metal properties") **INSERT BROWNSIDEBAR3** For example Quantum Molecular Dynamics (QMD) simulations currently provide the foundation

for many EOS tables. Present-day computing power enables the application of QMD to simple materials such as some insulators (e.g. carbon, both single- and multi-phase versions) and to simple metals (e.g. aluminum). Moving to more complex metals (an actinide like plutonium for example) increases the computing requirements significantly due to the much higher complexity associated with accurately describing the particle interactions within an actinide. Establishing a single actinide data point is just within reach today and requires 100's of teraflops computing. Producing a full data-table, allowing for the complex multiphase behavior of plutonium for example (sidebar: "Predicting the Fundamental Properties of Complex Materials") **INSERT BROWNSIDEBAR4** will require around 1000 data points and exascale computing capabilities.

More advanced methods for producing EOS data points are also being developed that will help quantify the uncertainties inherent in the QMD simulations by removing some of the fundamental approximations. These methods are significantly more expensive than QMD-based approaches and a recent estimation of the cost of producing a single multiphase actinide *data point* is found to be in the exascale regime.

EOS is not an isolated material data requirement for exascale computing. Similar examples, all of which demand exascale computing, could be given for most of the other material data areas of interest to the NNSA. Material strength, for example, is related to the ability of a material to deform. Deformation of crystalline material is dependent upon the behavior of an extremely large number of dislocations. A dislocation is essentially an irregularity in the crystal structure of a material, and their presence can influence many properties of a material. Modeling these dislocations and their evolution and growth can be performed at the atomic scale but very large numbers need to be modeled to determine their influence at the macroscopic level, as is needed for nuclear weapon simulations. Codes to track these dislocations have now been written and simulations that predict the gross material behavior, *under a single set of initial conditions*, for a single crystal of materials like tantalum, have just become possible with 100 teraflop class computing. To model some of the complex alloys that are present in a nuclear weapon will take much greater spatial coverage in the dislocation dynamics simulations, in order to accurately capture the effects of a statistically significant number of particulates, resulting in a ~10,000 increase in computing requirements, again leading to the need for exascale computing.

Uncertainty Quantification. The final example of a class of SSP problems that require exascale computing relates to the relatively new field of Uncertainty Quantification (UQ) for computational simulations. Gaining confidence in simulations, or assessing accuracy via validation, is not new, and is indeed an essential element of assessment via simulation. However as part of the drive to develop a predictive simulation capability, the NNSA has been developing methodologies to assess the numeric fidelity by quantifying uncertainty in its' predictions. The NNSA is not unique in this desire. Many researchers, in many different disciplines, turn to numerical

simulations to predict, or explain, complex phenomena; quantifying their simulation accuracy is of growing importance.

At present there is not an agreed methodology for quantifying simulation uncertainties. This is still a topic of research and the NNSA is assessing several options. However a key component in most of these methods is the need to measure the variability of the simulation results. In the limit where the simulation is solving the right problem, with the right equations, this variability will be strongly correlated to the simulation uncertainty.

For national security applications, the effect of small uncertainties in input variables can have a large effect on output results. This has led researchers to use a direct approach where an ensemble of calculations is run. This approach has been adopted by scientific, financial and industrial communities and applied to problems as diverse as the economy of the European Union and the burn in an inertial confinement fusion capsule. Naively, the direct approach would take an individual simulation which requires say, 20 parameters, and then specify the range of values that each of the 20 parameters could take. A good example would be to take the maximum, minimum and mid-point values of each parameter. With such an approach, the simulation would be performed for each of the different combinations of the parameter values leading to 3^{20} or about 9 billion simulations. In reality nuclear weapon simulations typically have hundreds, if not thousands, of potential parameters, and as noted above thousands of data points in their material property databases. The brute force approach is therefore totally impractical, even with exascale computing. This difficulty is called the "curse of dimensionality".

The UQ community has developed a variety of tools to handle this problem. It includes screening methods that seek to reduce the overall dimensionality of the space by reducing the number of input variables considered. Once the dimensionality has been reduced, statistical methods based on various sampling methodologies have been adopted. These statistical methods allow the researcher to sample from the full input space and develop a response surface that represents the output. A simple schematic example for a 2 dimensional parameter space is shown in the figure.

INSERT BROWNFIG1

In this manner the number of simulations required can be reduced to a manageable level. The NNSA has recently performed several such assessments, with 1D and 2D simulations, which have generally given meaningful results with about a thousand simulations. In the future, the goal of SSP predictive simulation is to be able to routinely run 3D predictive UQ studies. As noted above, a 3D nuclear weapon simulation, using the current established spatial resolution, requires ~100 teraflops. A sensitivity analysis in 3D will therefore require ~1000 times this degree of computing, i.e. 100's petaflops. Evidence from all the efforts so far imply that removing the UGT calibrations from SSP simulations will noticeably increase the

basic cost of a 3D simulation, so that 3D predictive UQ studies will be an exascale computing challenge.

Computational Support for Broader National Security Issues

Stockpile stewardship is not the only national security challenge that will require exascale computing in the coming decade. Exascale computing capabilities and tools, many leveraged from the ASC Program, can provide support for national security concerns such as nuclear forensics, nuclear counterterrorism, seismic modeling for non-proliferation, radiation hardening and survivability for microelectronics, vulnerabilities of critical infrastructure, cyber security (sidebar: "Cyber security analysis – a rapidly emerging driver for exascale computing") **INSERT BROWNSIDEBAR5** weapon effects and foreign assessments. Safety analysis, both for weapons systems and warfighter protection systems can also be supported by high performance computing (sidebar: "Safety analysis through Multi-Physics HPC Simulation"). **INSERT BROWNSIDEBAR6** In the past, ASC resources have also been used in the space shuttle Columbia investigation and in missile defense simulations.

Nuclear Counterterrorism. The nuclear threat has posed two main challenges, namely the need to eliminate the threat of a nuclear attack, and the minimization of the impact of such an attack should one occur. Furthermore, the nuclear threat has evolved in recent years. While the development of the U.S. stockpile was originally in response to the existence of large hostile nation-states, the concern now is with proliferation of nuclear technology to other nation-states as well as to non-state actors (e.g. terrorist organizations). A nuclear counterterrorism program must be able to detect the existence of nuclear weapons elsewhere and to determine the source of nuclear materials ("attribution"). Assessing these threats, evaluating weapons disablement technologies, and developing a capability to rapidly determine a nuclear weapons design from post-detonation debris should an attack occur, are all essential elements of such a program.

Computational technology needs for nuclear counterterrorism are potentially even greater than for stockpile stewardship. Improvised designs are expected to be much larger, physically, than weapons in the stockpile, and so analysis of these can be significantly more computationally intensive as they will require larger computational meshes to simulate. In addition, they often contain artifacts that necessitate expensive 3D modeling and the diversity of potential designs is extensive. If we are to be proactive in analyzing the risk posed by all likely potential designs, thousands of simulations will be required. The 3D calculations needed are not quite possible today, as computing achieves the petascale level. Running thousands of such simulations will therefore take exascale computing.

Assessing disablement technology effectiveness will also require large scale computing. For example, a single simulation of a shape charge jet using a realistic target will require an estimated 5 petaflop-days. In the case where the threat

knowledge is incomplete, hundreds of such calculations would be required to study the necessary potential weapons designs and disablement technologies, making this an exascale challenge. Exascale computing also opens up the possibility of performing these analyses in real time. Providing it is possible to use exascale computer performance to reduce the wall clock turnaround time of the simulation, a 5 petaflop-day simulation result could be available in a few minutes, potentially changing the role of simulation in this scenario.

In the event of a nuclear detonation, there will be tremendous pressure to provide as much information as quickly as possible on the type of device detonated to help enable identification of the perpetrators. This involves a backward engineering exercise where post-detonation diagnostics are matched to possible designs, again requiring numerous simulations. Using an example time frame of three days for the turn-around of such an analysis, one can estimate that the computing capacity required for future assessments, using 3D codes running suites of problems, to be in the exascale computing regime.

Nuclear non-proliferation is another area where HPC simulation technology will be essential. It is necessary to be able to determine existing and emerging weapons designs of our adversaries, locate and quantify existing nuclear materials, and evaluate the proliferation impact of the 'civil' nuclear renaissance (i.e. the development of nuclear power production technology in hostile states) based on intelligence information that is collected. It must also be possible to differentiate nuclear tests and other man-made explosions from naturally occurring events such as earthquakes in order to identify possible nuclear tests in other parts of the world. Computational seismology is an essential technology that supports this determination. Modern simulation tools and high performance computing allow for the modeling of explosion-generated seismic waves from the detonation underground, through the 3D earth to the observing seismic station. Current terascale computations are routine and petascale calculations have been demonstrated. However, to model seismic wave generation and propagation with sufficient fidelity for modern national security requirements, the development of code coupling and exascale computational power will be required. (sidebar: "Detonation-to- Detector, Simulation for Nuclear Explosion Monitoring").
INSERT BROWNSIDEBAR7

Summary

Computational simulation has been, and will continue to be, an essential element of U.S. national security programs. The NNSA ASC program has been driven since its inception by the need to ensure the safety, reliability and performance of the nation's nuclear weapons stockpile without nuclear testing. Even while the nation is not updating the nuclear stockpile, alterations are needed due to other factors e.g. aging. The challenge of underwriting these changes, without nuclear testing and with designers who increasingly do not have nuclear test experience, will be met by

enhancing the role of simulation so that it becomes a quantified predictive capability. This will be achieved by the development of high-fidelity, three-dimensional simulation codes, and the creation and deployment of the required computational capabilities and supporting infrastructure. With these capabilities computational science will be on an equal footing with theoretical and experimental science as a tool for studying basic issues of weapons science and for scientific discovery. Today, significant advances have been made in the scientific base of ASC simulation capabilities, but improvements are still required in the geometric fidelity, physics fidelity and numeric fidelity of national security simulations.

Three stockpile stewardship examples have been presented which represent the many different computational demands from the SSP. These were high resolution 3D simulations, development of material property data via simulation, and an assessment of the simulation uncertainty. Additionally it was outlined how these improved SSP capabilities are now used for broader national security missions and a number of specific examples were presented. In all cases mission requirements drive the need for a quantified predictive capability, which demands 3D simulations, either at significantly improved spatial and physics resolution, or in suites of very large numbers. To delivery these simulations, and therefore to meet the underlying mission requirements, requires exascale computing resources.

Sidebar1

History of ASC Program Computing Capability

The *Accelerated Strategic Computing Initiative* (ASCI) was initiated in 1995 as part of the DOE National Nuclear Security Administration's (NNSA) Science-Based Stockpile Stewardship Program (SSP) to develop the computational science and the computational horsepower needed to support the nuclear weapons mission. The essence of science-based stockpile stewardship is to replace underground nuclear tests (UGTs) with improved science and simulation, supplemented by specialized non-nuclear experiments.

The original ten-year ASCI goal was to develop and deliver computing platforms capable of sustained 100 teraFLOPS computing, along with improved physics-based models and three-dimensional multi-physics codes capable of running on them. This was an immense challenge and required effective teaming between Lawrence Livermore (LLNL), Los Alamos (LANL) and Sandia National Laboratories (SNL), the NNSA "tri-labs", the high-performance computing industry, and academia. The architecture pursued was massively parallel computing clusters with high-speed interconnects. The first significant delivery was the Intel based "ASCI Red" machine sited at SNL, the first system capable of one sustained teraFLOPS performance. Two years later, three teraflops systems, the SGI based "ASCI Blue Mountain" and IBM based "ASC Blue Pacific," were sited at LANL and LLNL, respectively. In 2000, the 12 teraFLOPs "ASCI White" machine was delivered to LLNL, representing an order-of-magnitude increase in computing power over the first five years of the initiative.

Steady progress continued to be made in the second half of the decade of ASCI, producing the Compaq based "ASCI Q" machine (20 teraFLOPS) at LANL and the Cray based "ASCI Red Storm" machine (40 teraFLOPS) at SNL in 2004. The original 100 teraFLOPS goal of ASCI was achieved in 2005 with the delivery of the "Purple" machine at LLNL in 2005. The capability of this machine represented a two-order of magnitude increase in computing capability over the first "ASCI Red" machine, and was 1000 times more powerful than computers prior to the advent of ASCI. Concurrent with these advances in computing hardware, ASCI also developed and deployed significant three-dimensional simulation capabilities, augmenting and replacing the "legacy" one-dimensional and two-dimensional simulation codes that had been used in the past.

The decade of ASCI demonstrated the ability to build computers and the associated infrastructure and environments capable of 100 teraFLOPS performance with sufficient reliability and availability to execute simulations that may run for days or weeks. Successful use of these leadership-class machines, along with physics-based simulation codes, has legitimized the science-based SSP strategy of replacing

nuclear tests with science and simulation. At the successful end of the decade long ASCI initiative, simulation had become a core element of the science-based SSP, and ASCI transitioned into the Advanced Simulation and Computing Program (ASC). In the absence of underground nuclear testing, ASC has become the *integrating element* of the SSP.

Progress in ASC computing capability has continued. At the same time as the Purple machine, ASC acquired the IBM based BlueGene/L system (360 teraFLOPS) whose development was jointly supported from the NNSA and DOE Office of Science. BlueGene/L was ranked as the most powerful computer in the world for three and a half years, and was used for cutting-edge work in simulation science resulting in three Gordon Bell prizes in supercomputing.

ASC's leadership-class computers have dominated the "Top500 list" of supercomputers over the last decade, but it has also had very significant impacts on the field of high-performance computing overall. Looking at data from a recent "Top500 list," the first 12 of the worlds most powerful computers directly benefitted from ASC funded architectures; 25% of the architectures of the top 500 computers are supported heavily by the ASC InfiniBand PathForward project; and 38% of the supercomputers on the "Top500 list" have benefitted heavily from ASC investments.

INSERT BROWNFIG2.PPT

Sidebar 2:

How a computational simulation is created

Advanced simulation techniques lie at the heart of many of the nation's most pressing scientific challenges, including understanding our changing climate, designing safe and efficient energy sources, and managing the nation's nuclear stockpile. For example, when assessing changes to the nuclear stockpile the impact of any such changes can be evaluated using computer simulations before the optimum design changes are chosen. These simulations are built using advanced mathematical models that describe the underlying physical phenomena, and sophisticated software tools that allow scientists to examine solutions for many different scenarios.

To build the simulations, research scientists first devise a mathematical model of the physical process they would like to study. For national security applications, this typically results in a system of many equations that approximate physical processes, along with a description of what is occurring on the boundary (boundary conditions) and at the beginning of the simulation (initial conditions). In addition, scientists must develop a computer representation of the computational domain. The geometry of the domain can be as simple as a rectangular box or sphere or as complex as one can imagine when studying advanced scientific devices. In most cases, the mathematical equations describing the physical phenomenon cannot be solved analytically on the domain of interest. Instead, the domain is decomposed into a collection of simpler geometries—a mesh—typically comprising triangles or quadrilaterals in two dimensions and tetrahedrons or hexahedrons in three dimensions. Once the mesh has been generated, the mathematical equations are approximated on that mesh resulting in a system of algebraic equations that is easier to solve on a computer than is the original equation.

As discussed in the main text, quantifying the uncertainty in a simulation is an important task. In addition to creating the basic numerical approximation to the physical equations, numerical parameters are often introduced, e.g. convergence and accuracy cut-offs. Many of these parameters cannot be established universally, depend upon the problem being solved and are therefore left as parameters that can be adjusted by the scientist running the code should they need to. Additionally the simulation will need to represent how a material behaves in response to various stimuli. For example “How quickly will a material heat up when a heat source is placed next to it”. Material properties like this can be given to the simulation in either tabular or analytic form. In either case, tabular values or parameter setting for the analytic forms need to be provided. Once again all these values are approximations to reality and may be changed for a simulation.

Having established a system of equations and the appropriate parameter and data settings, the simulation is extensively analyzed and, when possible, validated against experiments to ensure the solution is sufficiently accurate. A key influence on the accuracy of a numerical model, is the resolution of the computation, which is determined by the size of the computational meshes (the spatial resolution) and the size of the timestep (the temporal resolution). A property of a good numeric scheme is that as the mesh size and time step are reduced, the numerical approximation converges to the true solution. As a practical matter, simulations cannot be run with arbitrarily fine resolution since for a 3D, time-dependent problem, each factor of 10 increase in resolution incurs a factor of 10,000 increase in computational cost (see discussion in main text). Because of this implied limitation in possible resolution, the scientist must recognize that the numeric schemes cannot resolve phenomena that occur on length or time scales that are comparable to, or smaller than, the mesh size or time step, respectively. The validation process is therefore repeated with adjustments made to the mathematical model, the model parameters, the computational domain, the mesh, or the numerical solution process until either the optimum solution is found, if the solution is known, or the variability of the simulation is established.

INSERT BROWNFIG4.DOCX

Adapted from "ITAPS: Advanced Simulation Technologies for Application Scientists", SciDAC Review, Issue 13, Summer 2009 with permission of the author.

Sidebar 3

The role of computational resolution in determining transition metal properties[1]

Predicting material properties of transition metals. Predicting the performance of aging nuclear weapons requires detailed information on a material's phase transitions. For example, scientists want to know what happens to a metal as it changes from molten liquid to a solid and how that transition affects the material's characteristics, such as its strength. Because experiments at high pressures and temperatures are often difficult or impossible to conduct, scientists rely on computer models that have been validated with obtainable data. A material of interest to the stockpile stewardship program is the transition metal tantulum; computing its properties is very challenging computationally. In transition metals, the valence electrons, which interact with other elements to form compounds, are present in more than one shell. Thus, as tantalum solidifies, complex bonding structures form, and the transition from a melt phase to a solid can happen very slowly (from an atomic perspective). These physical processes are challenging to model. A molecular dynamics (MD) simulation of tantalum solidification may require billions of atoms, and the code must run many millions of time steps even though the process being simulated may last no more than a few nanoseconds. Researchers have been modeling systems with billions of atoms for a little over a decade. However, these models rely mainly on pair-potential techniques to describe the force each atom exerts on every other atom. Pair-potential techniques are effective for simple systems, such as those involving noble gases. Because noble gases have closed shells of electrons, the forces exerted on the atom are radially symmetric, resulting in spherically symmetric bonds. Pair-potential techniques do not model complex systems with the accuracy needed for stockpile stewardship research. Most of the transition metals—including tantalum—contain a partially filled *d*-band of electrons, which results in a more complicated bonding structure. For example, forces exerted on atoms are angularly dependent, and bonds may form between three or four atoms in a surrounding area. Accurately modeling these forces requires a sophisticated interaction potential. In 1990, NNSA physicist John Moriarty developed the model-generalized pseudo-potential theory (MGPT), which can be used to derive more accurate quantum-based interaction potentials. MGPT potentials are based on many-body expansions of a quantum-mechanically derived energy surface and include terms for two-, three-, and four-atom bonds. These potentials are validated by comparing information obtained by first-principles calculations and experiments.

---

[1] Material in this section is adapted from an article appearing in Science & Technology Review, Lawrence Livermore National laboratory, July/August 2006, and is used here with the permission of LLNL and the author.

Size matters. MGPT potentials are implemented in the ddcMD ("domain decomposition molecular-dynamics") code. With the ddcMD code running on BlueGene/L, weapons scientists can simulate billions of atoms on the necessary time scales to obtain reliable results. To study molten tantulum, a team of scientists from NNSA and IBM, modeled systems ranging from 64,000 to 128 million atoms compressed to 250 gigapascals of pressure at 5,000 kelvins. By varying the size of a simulation, the researchers gained confidence that their results were not affected by the size of the system being modeled. The 64,000-atom simulation showed two large grains, with a grain boundary spanning the simulation cell—an unrealistic result that indicated the system size was too small. In the simulation with more than 2 million atoms, the distribution of grain sizes was much more realistic. When the model size reached 16 million atoms, grain formation and growth were completely independent of system size. These simulations, for which the team won a 2005 Gordon Bell prize, are the first step toward modeling nucleation—when the transition to the solid phase begins—and growth in a manner that allows scientists to directly link processes at atomistic scales to those at micrometer scales and above. This early exploratory work, using classical MD methods, required the full capability of the ~0.5 petaflop BlueGene/L system, the most powerful supercomputer in the world at the time. These results indicate that more complex interaction potentials and simulation methods will be required in order to have confidence in the results. This will add orders of magnitude of complexity to these simulations, and will require exascale computing resources to predictively model this problem.

INSERT BROWNFIG4

Sidebar 4

Predicting the Fundamental Properties of Complex Materials[2]

In order to simulate the performance of weapons, one must be able to describe the mechanical, thermal, and nuclear behavior of the component materials throughout an extreme range of conditions. Nuclear weapons are constructed from a variety of unique materials, some of which are radioactive, and all of which change their behavior over time (age). In many cases it is not possible to experimentally measure material properties throughout the entire range of conditions to which they are subjected, nor is it always possible to directly assess the effects of aging. It is only through simulations of fundamental physics, supplemented by selected experiments, that satisfactory physics based models of material behavior can be constructed.

For example, plutonium is arguably one of the most complex materials known, and it is one of the least well understood. Before it liquefies, plutonium exhibits six solid material phases that vary considerably in density. Plus, a seventh phase may appear when the radioactive metal is under pressure.

To understand material phases, think of carbon and its most common solid phases: soft graphite and hard diamond. Both are made of carbon atoms, but the bonds that form between the atoms create two very different materials. Many elements have two or more solid phases, but most have no more than four. With six phases, solid plutonium is highly unusual. The material's peculiarities do not stop there. Experiments over the years have demonstrated other anomalous properties, including an almost complete absence of magnetism and highly unusual resistivity.

In an effort to explain some of plutonium's strange behavior and better understand results from past experiments, NNSA scientists and international collaborators performed some of the most precise predictions yet of delta-phase plutonium. For these first principles simulations, the team combined density functional theory (DFT) and dynamical mean field theory (DMFT) to calculate plutonium's delta-phase electronic structure, specifically its lack of magnetic "susceptibility." While DFT is useful for explaining the energy and interactions of many electronic systems, it may break down for certain properties of strongly correlated systems. Previous research combining the two theories to simulate delta-phase plutonium could obtain only approximate solutions to the DMFT equations. Recent advances in the continuous-

[2] Material in this section is adapted from an article appearing in Science & Technology Review, Lawrence Livermore National laboratory, November 2008, and is used here with the permission of LLNL and the author.

time quantum Monte Carlo (QMC) methods—and leadership-class supercomputing resourcesallowed exact solutions for the first time.

The team's simulations predicted that at room temperature with delta-phase plutonium at its equilibrium volume, the *f*-electrons are delocalized. That is, they easily move about the lattice and are not associated with one particular atom. Under these conditions, the material's magnetic susceptibility is very low above 600 Kelvin and only slightly higher at lower temperatures.

A change in temperature does not affect the magnetism of conventional metals such as platinum or molybdenum. In contrast, most strongly correlated materials other than plutonium exhibit a distinct relationship between temperature and magnetic susceptibility. Their magnetic susceptibility is high at low temperatures and lower at high temperatures. Experiments demonstrate, however, that plutonium's magnetic susceptibility is unlike that of other strongly correlated materials. In fact, it behaves more like a conventional metal, although it exhibits slight temperature dependence.

If the plutonium volume is expanded—if the lattice is stretched so that the *f*-electrons are farther apart—the magnetic susceptibility of plutonium changes. The Livermore simulations showed that as the lattice expands, the *f*-electrons become heavier, or localized. That is, they are more associated with one particular atom and thus cannot easily hop through the lattice. The plutonium is then a more strongly correlated material. The transition from delocalized to localized behavior occurs at increasingly lower temperatures as the lattice volume continues to expand. With greater distance between the electrons, delta-phase plutonium begins to behave more like other strongly correlated materials. As temperature drops, the metal's magnetic susceptibility increases.

The team intends to tackle plutonium's alpha phase next. Predicting the behavior of alpha-phase plutonium will be more challenging than the delta-phase simulations. The smallest individual crystal in delta-phase plutonium contains one atom. In the alpha phase, 16 atoms make up the smallest crystal. Another challenge will be to explain the role that other materials play in stabilizing delta-phase plutonium. Prediction of other plutonium properties, such as complete equation of state (EOS) data, are even more demanding of computing resources. Using today's most powerful computers, generating EOS data for materials as complex as plutonium is only tractable with (simpler) classical methods. Performing such calculations with more realistic first principles methods will require exascale computing power. Only with the aid of powerful supercomputers can researchers answer plutonium's many riddles.

INSERT BROWNFIG5

INSERT BROWNFIG6

Sidebar 5

Cyber Security Analysis – a rapidly emerging driver for exascale computing

Threats to the nation's cyber infrastructure, i.e. the large networked information systems on which business, energy and defense infrastructures increasingly depend, have become a serious national security concern. Understanding and addressing these threats will require computational capabilities that rival or exceed those for simulation and modeling of the physical systems of importance to national security.

INSERT BROWNFIG7

While modern cyber security operators at the high end tend to work with terabytes or petabytes of data, increased needs for complex computations and continually expanding datasets are already pushing the cyber security field towards exascale computing. For example, a single month of cooperative enterprise perimeter data collected by one government agency is known to contain tens of millions of distinct IP addresses – flows between source and destination – representing communication between one party and another. A consolidated 2D matrix representation of the implied connectivity graphs for this data would have a few quadrillion elements for each value of interest (e.g. first packet arrival time, port used, total bytes transferred). This simple connectivity matrix would be sparse (a few billion entries). A richer cybersecurity flow representation would involve representation of additional information: intermediate hops within a flow (devices besides the source/destination pair with access to the flow), dynamic and static address translation (allowing one IP address to map to different computers at different times, or allowing one IP address to represent multiple computers at a given time), and the need for detailed supplemental data about activities within any or all of the computers or devices along the flow, including intermediate hops. Combining these factors yields a complex multidimensional database of exascale size.

While real-time and near-term forensic analysis is important, so is longer-term retrospective analysis. Retrospective analysis extends the complexity through not just the address space, but across time. As an example, data collected at one-second time intervals for a sophisticated eighteen month cyber attack against the enterprise mentioned above would introduce over 25 billion new edges and nearly a million time steps. For either deep forensic analysis or "real-time" (aggregated to 1 minute intervals) cyber analysis and response, even the most basic analytical computations on a dynamic graph of over a

quadrillion elements and a million time steps is intractable on today's leadership class machines. And again, this is only considering network flows themselves, and not the critical interplay of activities within the hosts involved in the communication.

Emerging trends are to consider predictive as well as retrospective approaches to cyber defense. An example of a predictive approach is to use modeling and simulation to determine the effect of a new cybersecurity remedy on an Internet that includes IPv6 (Internet Protocol v6) traffic. IPv6, described as the successor to IPv4, has several advantages, among them increasing the directly addressable nodes on the Internet to about 3.4×10^{38}. Even with a simplified discrete network model that instantiates 340 "virtual machines" per physical node, simulation of this system would require a supercomputer with 10^{36} nodes –well beyond the exascale range.

Advancing cybersecurity analysis for very large scale and collaborative networks, particularly if the goal is to move from a "catch and patch" approach to one that is more proactive and predictive, will require fundamental advances in our understanding of the structure, mechanics, and dynamics of complex cyber networks. Progress is hindered by the enormous scale of the problem and the lack of scalable computing resources to perform high fidelity modeling, simulation, and analysis. Both operators seeking to protect systems, and researchers seeking to understand solutions, face computational complexity barriers. The ability to model, simulate, and analyze these problems is an important scientific challenge for global security – and, given the scale of the problem, one that will require extreme computing resources.

Sidebar 6:

Safety analysis through multi-physics HPC simulation

Safety analysis is an important element of national security. It is essential to know that the systems we design containing high explosives and propellants are inherently safe should an accident occur. We must also assure that the protection systems we provide our warfighters are maximally effective against enemy insult. High performance multi-physics simulations can provide critical insights for addressing these questions.

One example of an NNSA simulation tool used for safety analysis is the ALE3D code, which enables simulation for problems that must couple physics across a wide range of length and time scales. Example applications include high-explosive modeling from detonation to thermal cook-off, modeling of composite armor, penetrator mechanics, conventional weapon safety and performance, metal forging, concrete fracture rebar response, understanding how to protect of our forces through analysis of urban canyon blast, and safety systems such as armored vehicles and personal armor designed to protect against traumatic brain injury.

The figure demonstrates one small example of how we're continuously improving the fidelity of the physics, and the size and complexity of the systems that can be modeled, thanks to the rapid increase in computing power provided through programs like ASC. The ultimate goal – much like that of the stockpile stewardship program, is to provide predictive simulation capabilities that can replace much of the empirical work currently done through expensive experimentation and testing.

The ALE3D code has grown in capabilities over the last 20+ years due to a combination of new algorithms and the ability to exploit the growth in computational power of the largest available supercomputers. Analysts can build models on a laptop or workstation, and then easily scale their problems up to use the largest supercomputers available in the DOE and DoD. Exascale resources are required to allow analysts to not only run problems that are intractable on current generations of computers, but also to run suites of hundreds, or thousands, of calculations for parameter studies and uncertainty quantification.

INSERT BROWNFIG8

Sidebar 7

Detonation-to-Detector Simulation for Nuclear Explosion Monitoring

Seismology provides the best and most timely data to detect, locate and characterize a weapons system test (e.g. identify as an explosion, discriminate from background earthquakes, estimate the yield). When a nuclear device is detonated underground it vaporizes the rock immediately surrounding the device, creates a cavity and sends a shock wave into the earth within a few tenths of a second. This shock wave excites seismic waves that can be observed several minutes afterward at large distances, depending on the explosive yield, propagation pathway through the earth and observation conditions (e.g. background noise at a station). For example a 1 kiloton fully coupled (buried) explosion in strong rock can result in a Richter magnitude equivalent 4 event and be observed at distances of 1000's km. However, the level and character of ground shaking are strongly dependent on the geologic material in the immediate vicinity of the detonation and the path the waves travel through the earth from source to sensor. The earth is heterogeneous on all scales, that is its composition and physical properties vary in three-dimensions (3D) from the mineral grains that compose a single rock to the tectonic plates that form the outer layers of our planet. Not surprisingly seismic waves are strongly impacted by this 3D variability, especially in the near-surface crust and uppermost mantle (0-100 km depths). This makes every source-sensor path through the earth unique and introduces tremendous uncertainty in seismic modeling.

During the Cold War, nuclear explosion monitoring was focused on a few well-known test sites and empirical analysis of seismic recordings from past events. Now national security requirements are concerned with proliferation and emerging nuclear states and this requires broad area monitoring without reliance on past explosions. Equally challenging is the need to monitor some areas devoid of earthquakes for comparison. These requirements stimulate the development of simulation and modeling tools to gain a predictive capability for seismic waves emerging from explosions and/or earthquakes from any source location to any sensor, especially in areas of proliferation concern. Modern simulation tools and high performance computing allow for the modeling of explosion-generated seismic waves from the detonation underground, through the 3D earth to the observing seismic station (sensor). Current terascale applications are routine and petascale calculations have been demonstrated. However, in order to model seismic wave generation and propagation to sufficient fidelity across the broad range of spatial (10^{-2} – 10^{6} m) and temporal scales (10^{-6} – 10^{3} s) requires development of code coupling and giant advances in computational power.

Full-scale 3D end-to-end (device-to-detector) modeling is the goal for computational nuclear explosion seismology and will require exascale computing. This will enable seismologists to better detect weak signals from small distant events, provide improved estimates of event location and explosive yield as well as reduce uncertainty in source type discrimination. Toward this goal we are modeling elements of the full-scale problem and working to piece them together. Presently we can model shock wave generation in different rock types. Similar to the SBSS program, calibration of material models with past nuclear test data is essential to build confidence in simulation results for the shock waves generated by underground nuclear tests. To improve propagation models on continental scales, earthquake data can be used to image 3D earth structure. New methods based on fully 3D simulation of earthquake motions and time-reversal (adjoint waveform tomography) are being applied to resolve detailed sub-surface structure. While these methods have been demonstrated, they will require exascale computing to perform the large number of high-resolution calculations for 100's of earthquakes and iterations over model optimizations. The future is bright for exascale computing to enable major advances in seismic nuclear monitoring.

INSERT BROWNFIG9

INSERT BROWNFIG10

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

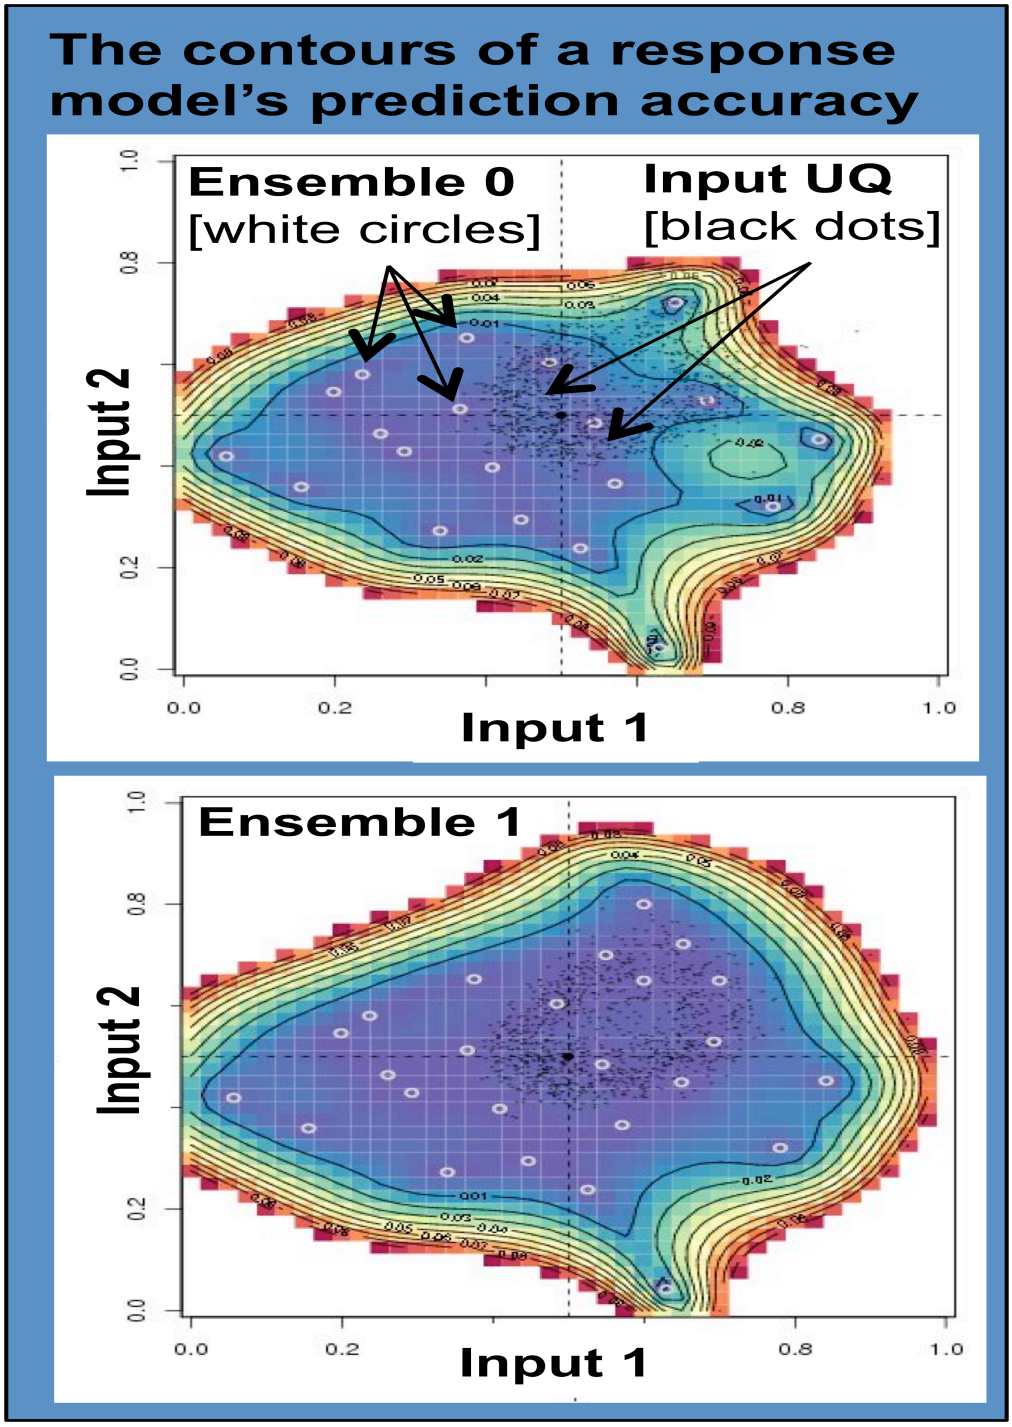

Fig 1: White circles in part A identify simulation tests run varying two parameters. Based on those answers, results can be predicted for nearby values of the two parameters and a predicted accuracy for those values can be calculated. In this case, purple is highest predicted accuracy. Based on this result, adaptive sampling recommends more samples in the upper right area, thus improving the overall accuracy shown in part B.

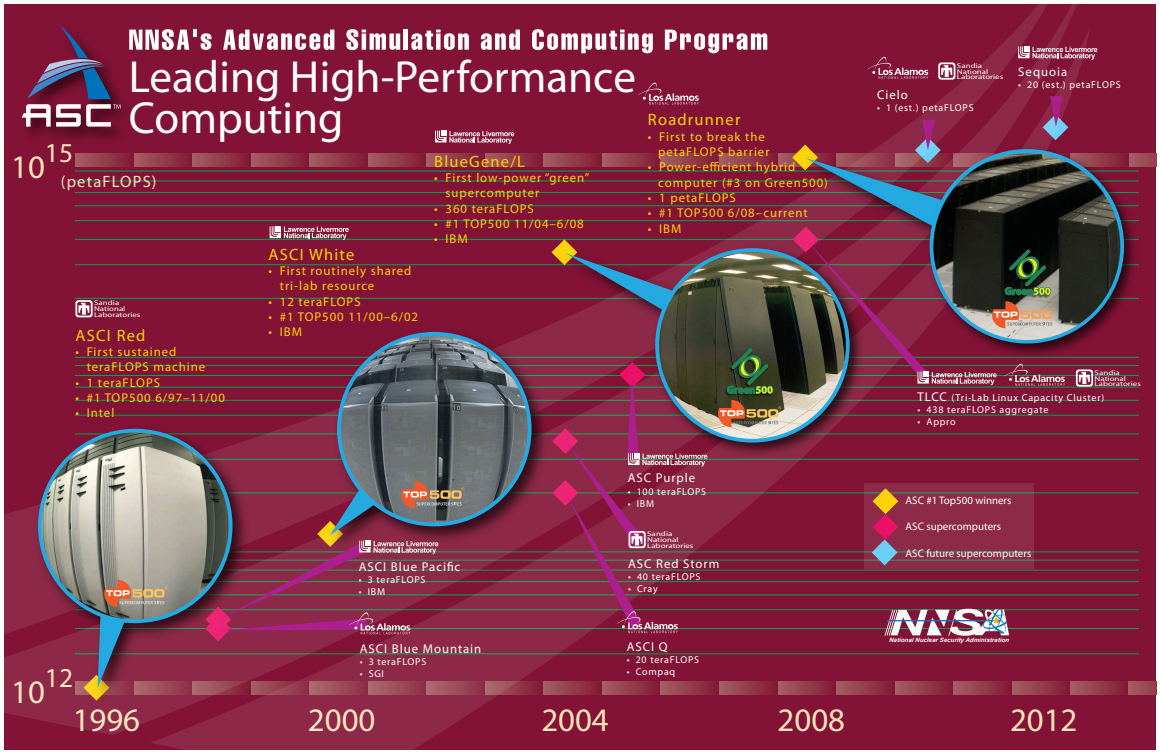

Fig 2 : The performance of supercomputers in the ASC program increased by three orders of magnitude in just ten years and is expected to grow another order of magnitude in the next two. Here, the peak performance of each of the ASC platforms is shown plotted against delivery date.

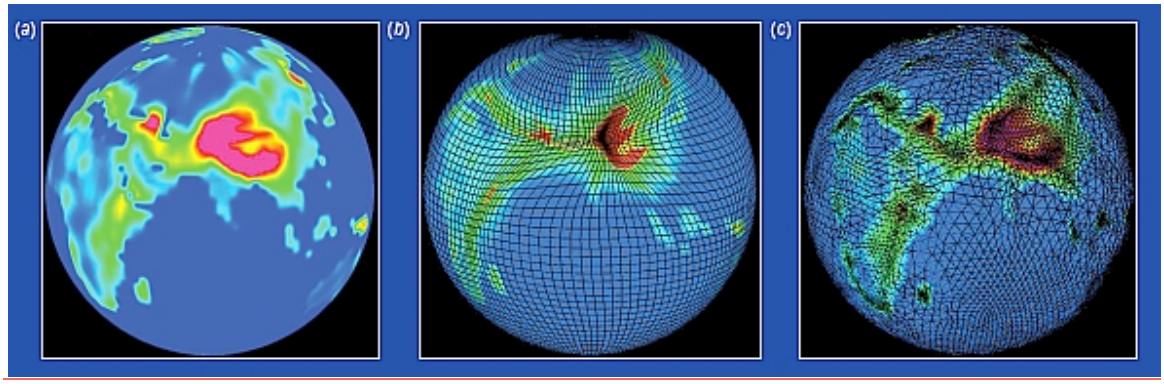

Fig 3: Different types of meshes can be used to represent geometry in a simulation. This figure represents (a) orography (height of the terrain) with color, and uses structured (b) and unstructured (c) meshes to concentrate mesh points in regions of interest.

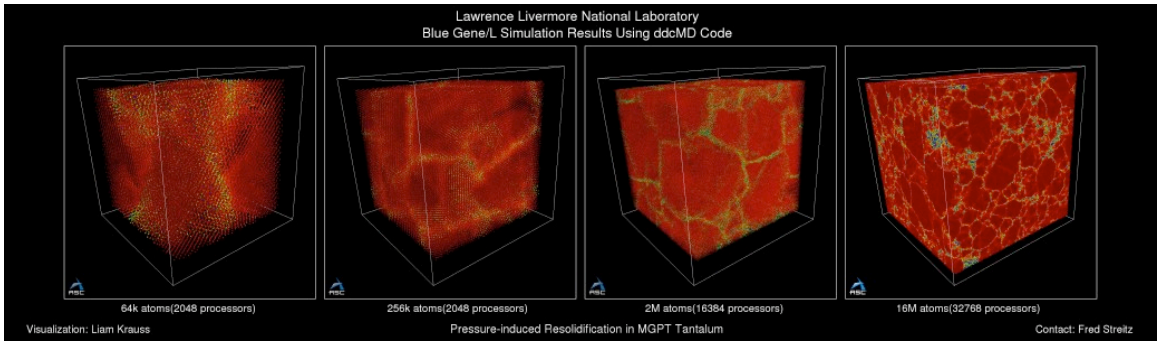

Fig 4: The NNSA-IBM team investigated the transition of tantalum from melt to solid as a function of system size. The results for (a) 64,000 atoms, (b) 256,000 atoms, (c) 2,048,000 atoms and (d) 16,384,000 atoms show that simulations smaller than about two million atoms produced artificial grain boundaries during the solidification (see cross-section in insets). The results from the 16 million-atom simulation exhibit the formation of realistic grains nucleating and growing directly from the melt.

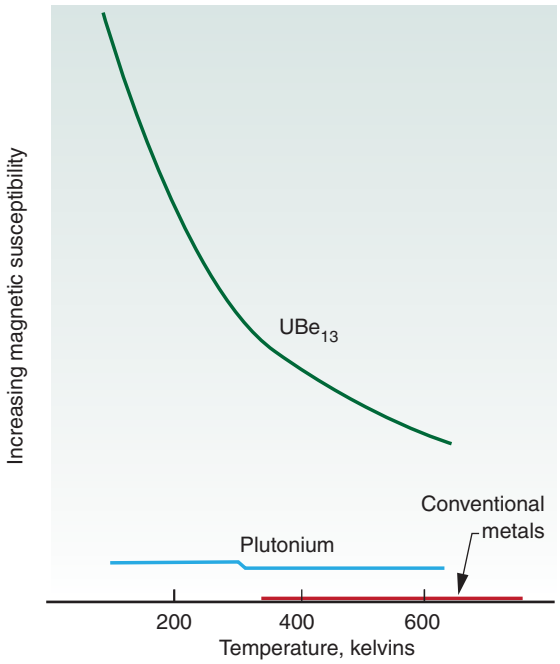

Fig 5: The magnetic susceptibility of plutonium is unusual. For conventional metals such as molybdenum, titanium, and platinum, magnetism does not change with temperature (red curve). Strongly correlated materials, however, are more magnetic at low temperatures than they are at higher temperatures. The green curve shown for uranium–beryllium-13 (UBe_{13}) is typical. Plutonium's magnetic susceptibility (blue curve) lies between these cases.

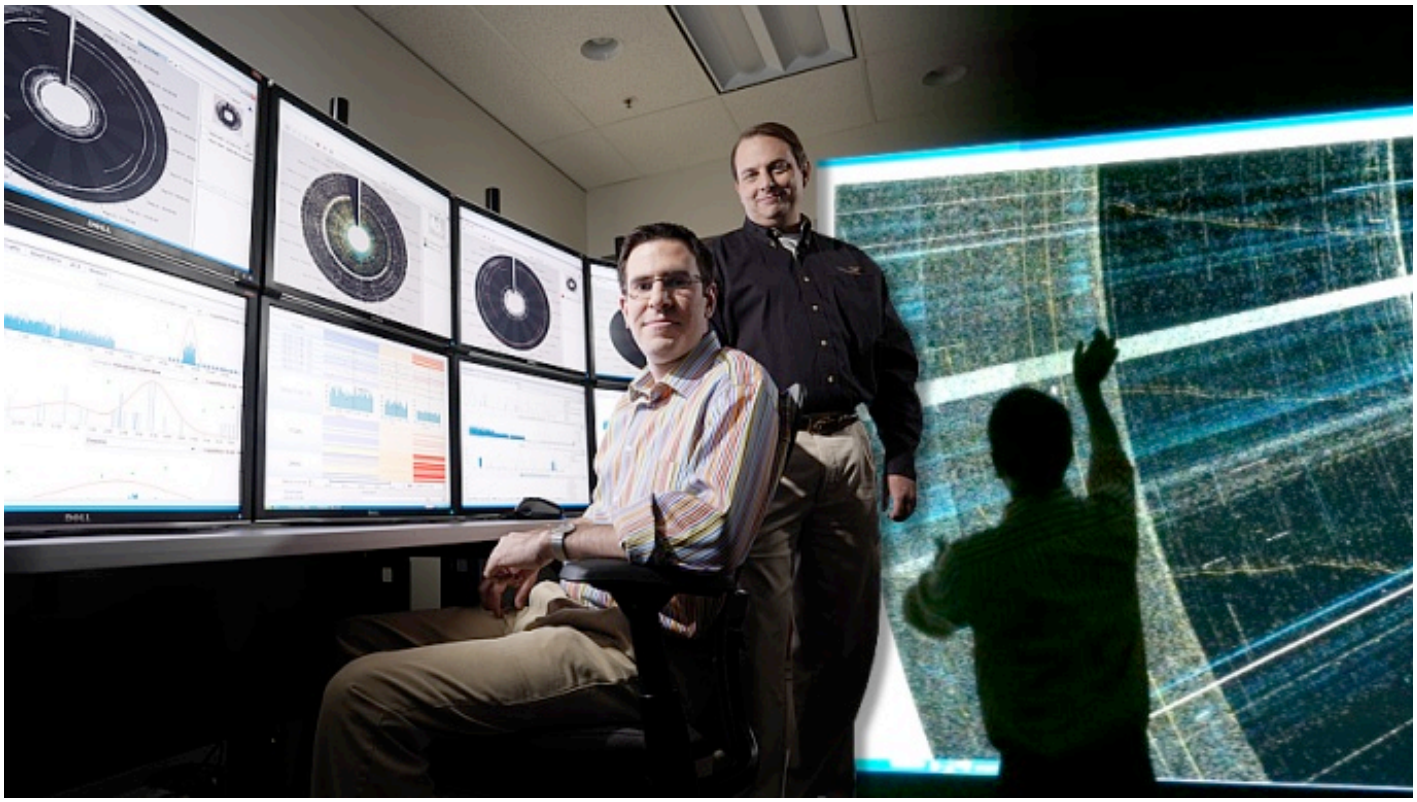

Fig 7: Bill Pike and Jeff Mauth analyzing network traffic with PNNL's Traffic Circle

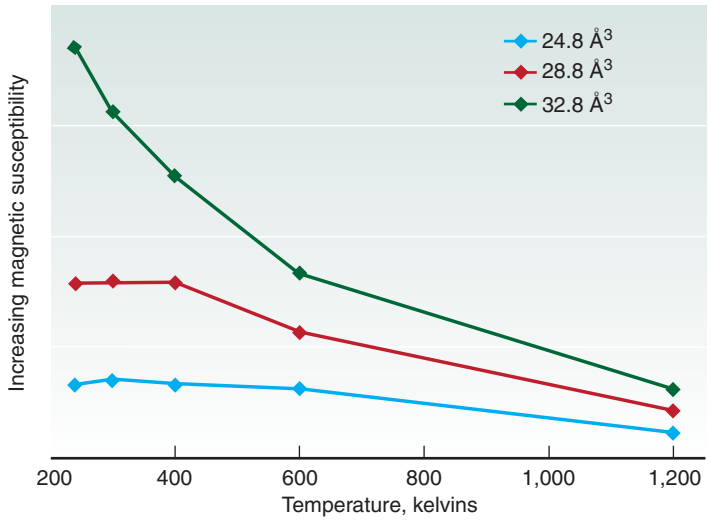

Fig 6: In the NNSA simulations, the volume of plutonium was increased, pulling the plutonium atoms farther apart. Only under these circumstances and at the largest volume does plutonium's magnetic susceptibility begin to mimic the temperature dependence of other strongly correlated materials such as UBe_{13}. (Å^3 = cubic angstroms, where 1 Å^3 equals 1×10^{-30} cubic meters.)

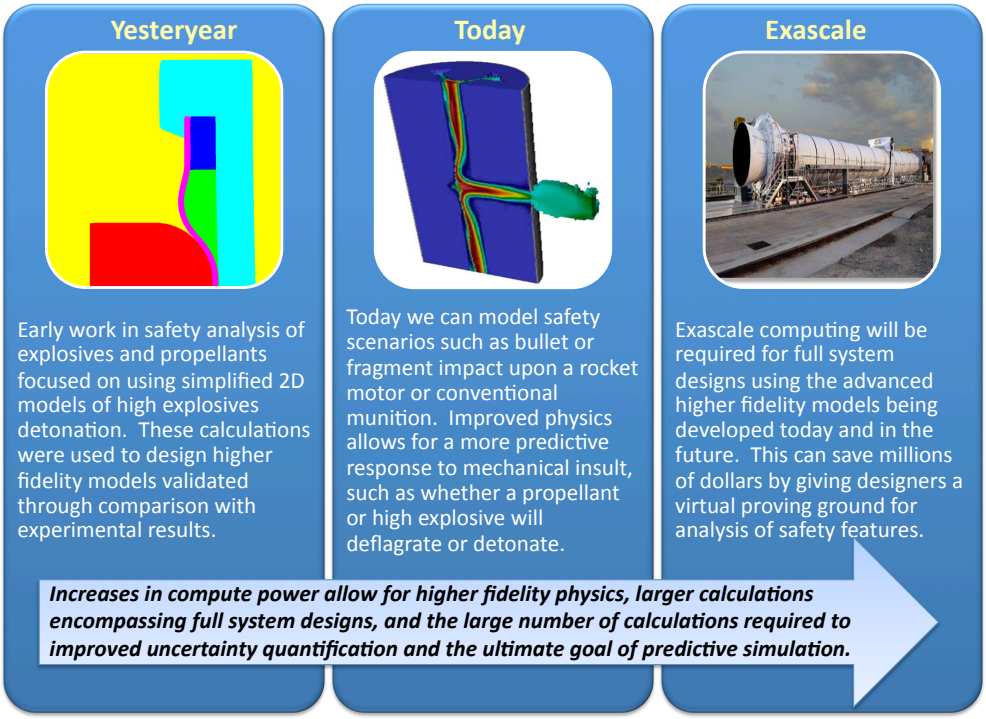

Fig 8: (no caption – info is included in graphic)

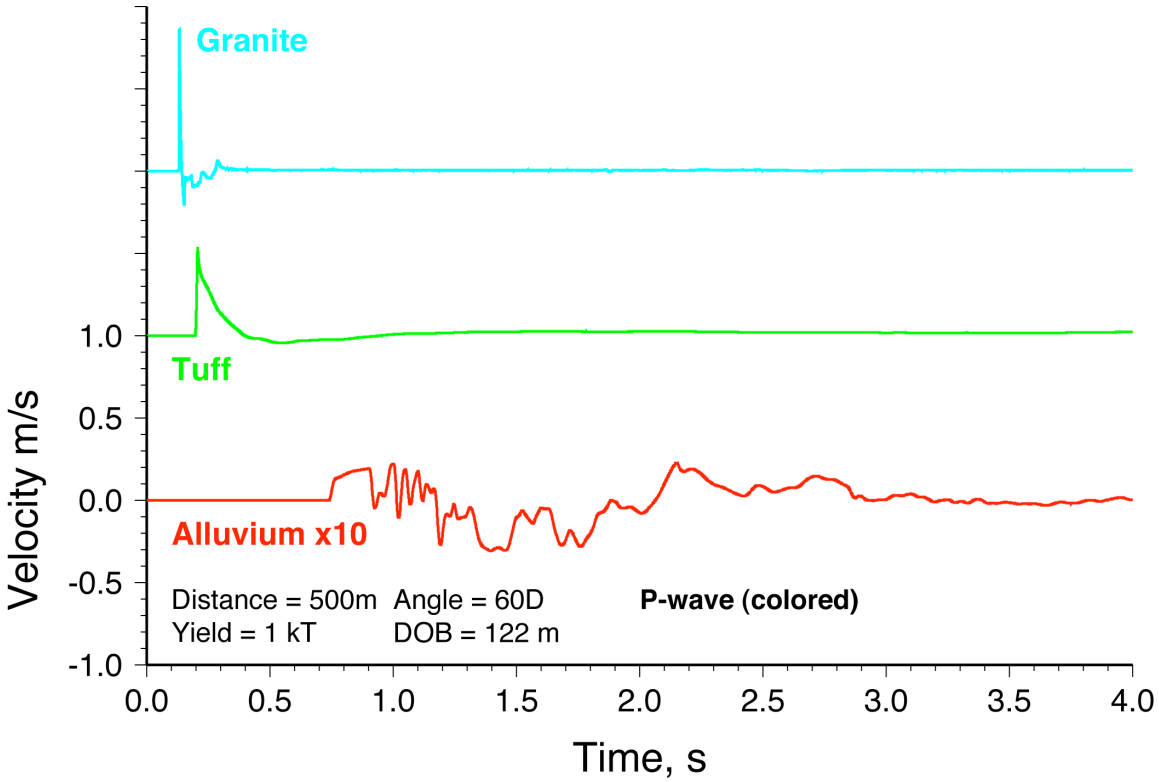

Fig 9: Modeling of explosion generated shock waves in different material models (granite, volcanic tuff and sedimentary alluvium) results in vastly different wave motions. Material models were generated to reproduce observations from a large set of legacy ground motion data from past nuclear tests.

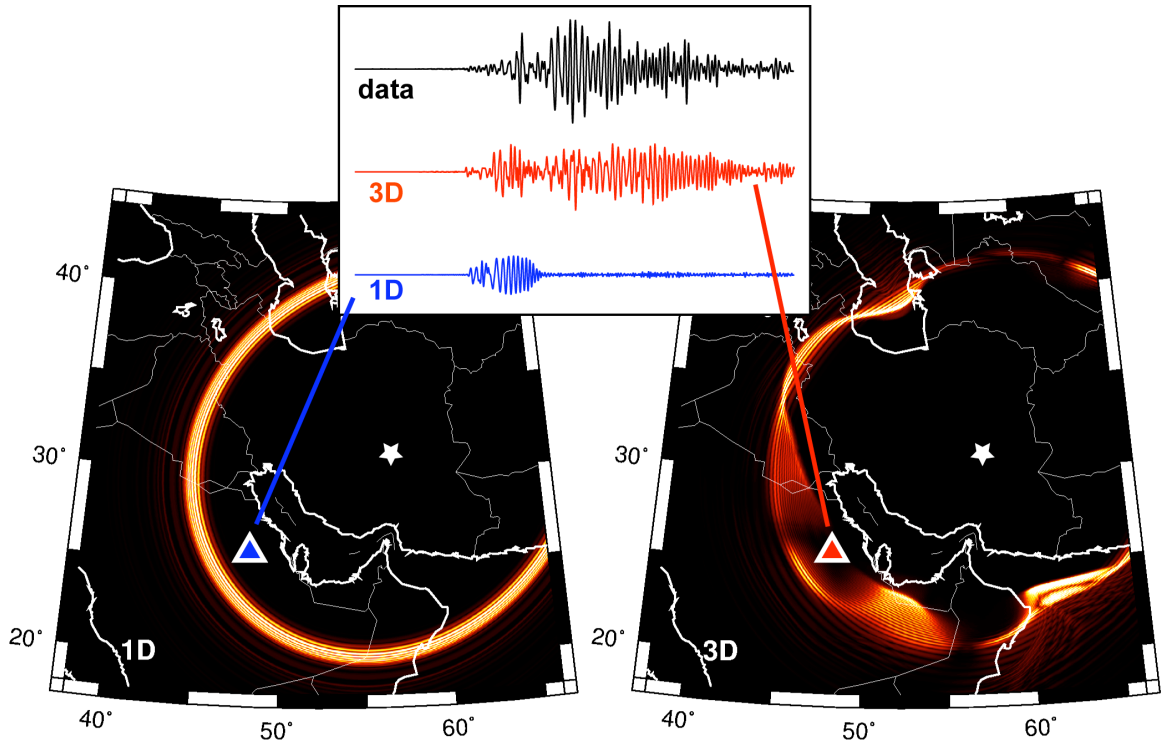

Fig 10: Simulations of a seismic event in the Middle East using an unrealistic one-dimensional (1D, left) and three-dimensional (3D, right) model. Comparison with an observed seismogram (top, black) shows that the 3D model (red) predicts the long-duration character of the data and the 1D model (blue) is not at all like the data.